

Climbing Depth-bounded Adjacent Discrepancy Search for Solving Hybrid Flow Shop Scheduling Problems with Multiprocessor Tasks

Asma LAHIMER¹, Pierre LOPEZ¹, and Mohamed HAOUARI²

¹ CNRS ; LAAS ; 7 avenue du colonel Roche, F-31077 Toulouse Cedex 4, France
Université de Toulouse ; UPS, INSA, INP, ISAE ; UT1, UTM, LAAS ; F-31077
Toulouse Cedex 4, France

`asma.lahimer@laas.fr, pierre.lopez@laas.fr`

² INSAT ; Institut National des Sciences Appliquées et de Technologie
Centre Urbain Nord BP 676 - 1080 Tunis Cedex, Tunisie
`mohamed.haouari@insat.rnu.tn`

Abstract. This paper considers multiprocessor task scheduling in a multistage hybrid flow-shop environment. The problem even in its simplest form is NP-hard in the strong sense. The great deal of interest for this problem, besides its theoretical complexity, is animated by needs of various manufacturing and computing systems. We propose a new approach based on limited discrepancy search to solve the problem. Our method is tested with reference to a proposed lower bound as well as the best-known solutions in literature. Computational results show that the developed approach is efficient in particular for large-size problems.

Keywords: Hybrid flow shop scheduling, Multiprocessor tasks, Discrepancy search

1 Introduction

Flow shop scheduling refers to a manufacturing facility in which all jobs visit the production machines in the same order. In hybrid flow shop scheduling, the jobs serially traverse stages following the same production route, and must be assigned to one of the parallel machines composing each stage. The hybrid flow shop scheduling problem with multiprocessor tasks is itself a generalization of the hybrid flow shop problem, allowing tasks to be processed on more than one processor in a given stage, at a time. It can also be viewed as a specific case of the resource-constrained project scheduling problem (RCPSP).

Many applications of hybrid scheduling problems with multiprocessor tasks can be found in various manufacturing systems (*e.g.*, work-force assignment in [6], transportation problem with recirculation in [4]), as well as in some computer systems (*e.g.*, real-time machine-vision [8]).

Hybrid flow shop scheduling problem with multiprocessor tasks has received considerable attention from researchers and has been solved by various approaches, *e.g.* genetic algorithms [14], tabu search, and ant colony system [19].

Motivated by the success of discrepancy search for solving shop scheduling problems, in particular hybrid flow shop [2], [3], we propose in this paper a new approach based on discrepancy search to solve the hybrid flow shop problem with multiprocessor tasks.

2 Problem Definition

The hybrid flow shop scheduling problem with multiprocessor tasks can be formally described as follows: A set $\mathbf{J}=\{1, 2, \dots, n\}$ of n jobs, have to be processed in m stages. Hence, a job is a sequence of m tasks (one task for each stage). Each stage $\mathbf{i} = \{1, 2, \dots, m\}$ consists of m_i identical parallel processors. In a stage i , the job j requires simultaneously $size_{ij}$ processors. That is, $size_{ij}$ processors selected at stage j are required for processing job j for a period of time equal to the processing time requirement of job j at stage i , namely p_{ij} . The objective is to minimize the *makespan* (C_{\max}), that is, the completion time of all tasks in the last stage. According to the classical 3-field notation in production scheduling, the problem is denoted by $Fm(m_1, \dots, m_m) | size_{ij} | C_{\max}$.

3 Discrepancy Search

3.1 General Statement

Limited discrepancy search (LDS) was introduced in 1995 by Harvey and Ginsberg [9]. This seminal method can be considered as an alternative to the branch-and-bound procedure, backtracking techniques, and iterative sampling. From an optimization view-point this technique is similar to variable neighbourhood search. Indeed, it starts from an initial global instantiation suggested by a given heuristic and successively explores branches with increasing discrepancies from it, in order to obtain a solution (in a satisfaction context), or a solution of better performance (in an optimization context). A discrepancy is associated with any decision point in a search tree where the choice goes against the heuristic. For convenience, in a tree-like representation the heuristic choices are associated with left branches while right branches are considered as discrepancies. Since LDS proposition in 1995, several variants were suggested, among them, Improved Limited Discrepancy Search (ILDS) [12], Depth-bounded Discrepancy Search (DDS) [21], Discrepancy-Bounded Depth First Search [1] and Climbing Discrepancy Search (CDS) [13].

In the following sections, we focus on those methods that inspired our approach, in particular DDS and CDS.

3.2 Depth-bounded Discrepancy Search

Depth-bounded Discrepancy Search (DDS) developed in [21], is an improved LDS that prioritizes discrepancies at the top of the tree to correct early mistakes

first. This assumption is ensured by means of an iteratively increasing bound on the tree depth. Discrepancies below this bound are prohibited. DDS starts from an initial solution. At i th iteration, it explores those solutions on which discrepancies occur at a depth not greater than i .

3.3 Climbing Discrepancy Search

Climbing Discrepancy Search (CDS) is a local search method adapted to combinatorial optimization problems proposed in [13]. CDS starts from an initial solution that would be dynamically updated. Indeed, it visits branches progressively until a better solution is reached. Then, the initial solution is updated and the exploration process is restarted.

4 Proposal: Climbing Depth-bounded Adjacent Discrepancy Search

4.1 CDADS: Main Features

To stick to the problem under consideration, we now consider an optimization context. We propose CDADS (Climbing Depth-bounded Adjacent Discrepancy Search) method, that is a combination of a depth-bounded discrepancy search and a climbing discrepancy search. We also assume that, if several discrepancies occur in the construction of a solution, these discrepancies are necessarily *adjacent* in the list of successive decisions. CDADS starts from an initial solution obtained by a given heuristic, and explores its neighborhood progressively, according to the depth-bounded discrepancy search strategy. Hence, a limit depth d is fixed. Discrepancies below this bound are prohibited. At the i th iteration, we allow i discrepancies above the limit level d .

When considering solutions with more than one discrepancy, we require these discrepancies are achieved consecutively, that means a solution consists of discrepancies that happen one after the other. This assumption of adjacency considerably limits the search space. We also consider that the initial solution is generated by a ‘good’ heuristic. Thus, only the immediate neighborhood of a discrepancy may receive an additional discrepancy. We obtain a truncated DDS based on adjacent discrepancies, DADS (Depth-bounded Adjacent Discrepancy Search). This approach is illustrated by an example on a binary tree of depth 3 (see Figure 1). At the starting point, DADS visits the initial solution recommended by the heuristic. For convenience, we assume that left branches follow the heuristic. At first iteration, DADS visits leaf nodes at the depth limit with exactly one discrepancy. The first line shown under the branches reports the visit order of considered solution, while the second line illustrates the number of discrepancies made in each solution. The 2nd iteration allows to exploring more solutions with two discrepancies with respect to the adjacency assumption. In this representation, the maximum depth bound is taken to be 3. If now, we limit

the depth to two levels, several branches would not be retained, namely the branches 4, 6, and 7 would not be visited by DADS.

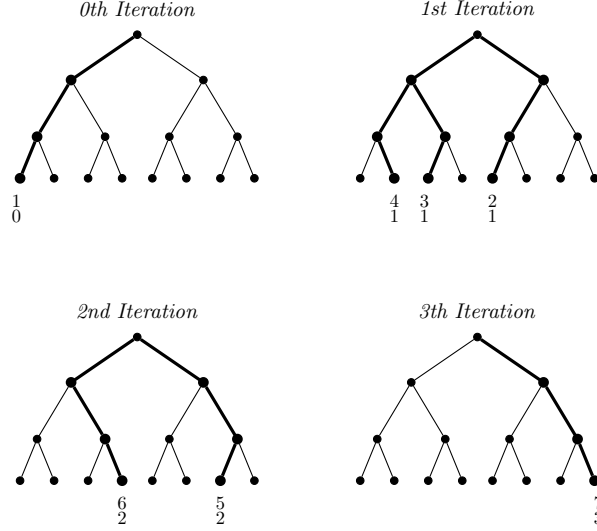


Fig. 1. Depth-bounded Adjacent Discrepancy Search

Going back to the optimization issue, CDADS merges the DADS strategy with a CDS exploration principle, that is the initial solution used by DADS is dynamically updated when a best solution is found, and the exploration process is restarted.

4.2 Heuristics

CDADS is strongly based on the quality of the initial solution. Thus, we carried out an experimental comparison between various priority rules presented in the literature [19], [15]. We considered the most effective heuristics to multiprocessor task hybrid flow shop scheduling. The four selected rules are:

- **SPT** (Shortest Processing Time), which ranks jobs according to the ascending order of their processing times;
- **SPR** (Shortest Processing Requirement), which ranks jobs according to the ascending order of their processing requirement;
- the **Energy** rule, considering first the jobs with the smallest energy (where the energy of an operation j at a stage i is evaluated by $p_{ij} \times size_{ij}$); and
- **NSPT_LastStage** (Normalized SPT applied at the last stage). For this latest rule, Şerifoğlu and Ulusoy [19] propose to schedule jobs according to their ranking index (RI_j) defined by:

$$RI_j = \frac{\max_k \{p_{mk}\} - p_{mj} + 1}{\max_k \{p_{mk}\} + 1}.$$

In Table 1, the selected priority rules are ranked according to their percentage of best solutions found, that is, performance.

Table 1. Heuristic selection

Priority Rule	Performance (%)
NSPT_LastStage	27
Energy	25
SPT	17
SPR	14

4.3 Schedule Generation Scheme

Schedule generation schemes (SGSs) are widely used in solving preemptive problems. We distinguish between serial SGS and parallel SGS. These two heuristics ensure task scheduling based on a given priority rule. Hence, tasks are selected one after the other and a start time is fixed for each one.

Serial SGSs are introduced in [11]. At each iteration, the first available task in ζ is selected, where ζ is the priority list recommended by the priority rule. The selected task is scheduled as soon as possible with respect to both resource constraints and precedence constraints.

Parallel SGSs developed in [5], suggest a chronological procedure in scheduling tasks. At each time t , a set ζ_t of tasks being scheduled is defined: this set contains unscheduled tasks that can be processed at t without breaking neither precedence constraints nor resource constraints. If we consider that \underline{t} is the first time where $\zeta_{\underline{t}} \neq \emptyset$, the first task in the priority list ζ belonging to $\zeta_{\underline{t}}$ is performed at \underline{t} . The same process is applied until all tasks are scheduled. The two schemes depicted above may appear similar. However, the schedule they generate are different: a serial SGS provides an active schedule while a parallel SGS generates a non-delay schedule.

In the scheduling theory, Sprecher *et al.* [20] show that the set of active schedules includes at least one optimal solution. On the contrary, non-delay schedules may eliminate all optima.

Concerning our method CDADS, we do not enumerate all possible solutions, so even serial SGSs may exclude all optimum solutions. Furthermore, in practice, parallel SGSs are known for their operational efficiency. Hence, we opt for the implementation of a parallel SGS which has been proved, moreover, to be more efficient in our experimental studies.

4.4 Lower Bound

For efficiency purpose, we join CDADS with an evaluation of lower bounds at each node. The proposed lower bound is based on lower bounds previously presented in [14]. Thus, we suggest this formula:

$$LB = \max(LB^s, LB^j)$$

where LB^j is a job-based lower bound similar to the one suggested in [14]:

$$LB^j = \max_{j \in J} \left(\sum_{i=1}^m p_{ij} \right); \text{ and } LB^s \text{ is a stage-based lower bound: } LB^s = \max_{i=1..m} LB(i).$$

For this latter bound, we claim that:

$$LB(i) = \begin{cases} \max[M_1(i), M_2(i), \max_{j \in J}(p_{ij})] + \min_{j \in J} \left(\sum_{l=i+1}^m p_{lj} \right), & \forall i = 1 \\ \min_{j \in J} \left(\sum_{l=1}^{i-1} p_{lj} \right) + \max[M_1(i), M_2(i), \max_{j \in J}(p_{ij})] + \min_{j \in J} \left(\sum_{l=i+1}^m p_{lj} \right), & \forall i = 2..m-1 \\ \min_{j \in J} \left(\sum_{l=1}^{i-1} p_{lj} \right) + \max[M_1(i), M_2(i), \max_{j \in J}(p_{ij})], & \forall i = m \end{cases}$$

where

$$M_1(i) = \left\lceil \frac{1}{m_i} \sum_{j \in J} (p_{ij} \text{size}_{ij}) \right\rceil$$

and

$$M_2(i) = \sum_{j \in A_i} p_{ij} + \frac{1}{2} \sum_{j \in B_i} p_{ij},$$

with

$$A_i = \{j | \text{size}_{ij} > \frac{m_i}{2}\}$$

and

$$B_i = \{j | \text{size}_{ij} = \frac{m_i}{2}\}.$$

Justification of the expression of $LB(i)$.

We assume that only non-delay task scheduling is considered.

The first term of $LB(i)$ gives a lower bound on the beginning of every job $j \in J$ on any machine of stage i .

The last term can be explained accordingly, since it is associated with the minimal required time to achieve the processing of every job j on all the subsequent stages of stage i .

The middle term concerns the processing of jobs on stage i . $M_1(i)$ stands for the mean stage load for job preemptive scheduling, while $M_2(i)$ reviews two different situations for partitioning the jobs according to their resource requirement. Set A_i consists of jobs that must be processed sequentially (resource requirement greater than the half of the resource capacity m_i). Set B_i groups together the jobs having a resource requirement exactly equal to the half of the resource capacity. Obviously, a job belonging to A_i and another job belonging to B_i must also be processed sequentially. The added term $\max_{j \in J}(p_{ij})$ contributes to maximize the evaluation of stage load on a considered stage i , especially when some jobs having high processing time are being scheduled.

This justifies the validity of the bound. \square

5 Computational Study

5.1 Test Beds

For comparison purpose, we assess the performance of CDADS on instances of Oğuz's benchmark available on her home page: http://home.ku.edu.tr/coguz/public_html/. This benchmark is widely used in the literature [18], [10], [16].

The number of jobs is taken to be $n = 5, 10, 20, 50, 100$ and the number of stages m takes its value from the set $\{2, 5, 8\}$. The benchmark considers two types of problems, "Type-1" and "Type-2". In 'Type-1' instances, the number of processors m_i available at each stage i (resource capacity) is randomly determined from the set $\{1, \dots, 5\}$, while in 'Type-2' m_i is fixed to 5 processors for every stage i . In fact, 'Type-2' instances are globally more flexible than 'Type-1 instances'. For each combination of n and m , and for each type, 10 instances are randomly generated, which leads a total of 300 instances. The processing time of each job j in stage i (p_{ij}) and its processing requirement ($size_{ij}$) are integer and are randomly generated from sets $\{1, \dots, 100\}$ and $\{1, \dots, m_i\}$, respectively.

The algorithm implementing CDADS was coded in C++ and run on an Intel core 2 Duo 2 GHz PC. The maximum CPU time is set to 60 seconds. The exploration is also stopped when CDADS reaches a given lower bound on the makespan. Obviously, if CDADS misses the optimal solution, the best-found solution when the maximum CPU time is reached, is then taken to be the problem solution.

5.2 Restart Policy

For the computational study, we have then retained four priority rules to generate the initial solutions (see Section 4.2). That is why we have introduced a restart policy to benefit from these heuristics. At a starting point, we use the best

rule, that is the NSPT_LastStage. However, if no improvement is noticed during the CDADS search, we restart the process with another solution obtained by applying the next rule “Energy” that could lead a more efficient solution for this specific instance, and so on.

The restart policy is limited by the size of the heuristics pool: restarts are then allowed at most four times, since we have selected four rules. At each restart k (starting from $k = 0$), we increase the number of maximum nodes that can be visited according to a geometrical series $nbrNodes \times f^k$, where f is fixed to 1.3 and $nbrNodes$ varies linearly with the problem size (the number of jobs n ; for example for $n = 20$ we fix $nbrNodes$ to 2000 nodes). Hence the search space is expanded at each restart.

5.3 Results

We tested two strategies for applying discrepancy: Top First and Bottom First. In the Top First exploration, discrepancies at the top of the tree are privileged while the Bottom First strategy favors discrepancies at the bottom. Computational study shows that CDADS is really more efficient with a Top First strategy (then contradicting – for the problem at hand – the statement of relative indifference of discrepancy order by [17]). Thus, the results shown below refer to this latter strategy.

Table 2 gives for each configuration (n : number of jobs, and m : number of stages) and each type, the average percentage deviation ($\%dev$) and the average CPU time. The average percentage deviation is measured in two ways:

- For small problems, solutions are compared to the optimal solutions (C_{\max}^* denotes the optimum makespan):

$$\frac{C_{\max} - C_{\max}^*}{C_{\max}^*} \times 100;$$

- For larger problems, solutions found by the CDADS are compared to the lower bound (LB):

$$\frac{C_{\max} - LB}{LB} \times 100.$$

As explained in Section 5.2, CDADS is run four times on each of the selected priority rules (NSPT_LastStage, Energy, SPT, SPR) for each instance. The best solution is taken to be the CDADS solution for the corresponding problem. According to findings of [19], the $Fm(m_1, \dots, m_m) | size_{ij} | C_{\max}$ problem and its symmetric have the same optimal makespan. Referring to this property, we apply a two-directional planning (forward schedule and backward schedule).

From Table 2, it is observed that the average percentage deviation is higher for ‘Type-1’ instances. Globally, $\%dev$ is 1.66% for ‘Type-1’ problems and 6.39% for ‘Type-2’ problems. This increase can be linked to several assumptions: the lower bound becomes less effective as m_i increases in ‘Type-2’ instances and so

Table 2. CDADS performance

n	m	‘Type-1’ Problems		‘Type-2’ Problems	
		$\%dev$	$CPU(s)$	$\%dev$	$CPU(s)$
5	2	0	< 0.1	0	< 0.1
	5	0.21	< 0.1	0.46	< 0.1
	8	1.71	< 0.1	0.5	< 0.1
10	2	0	< 0.1	1.72	< 0.1
	5	0.66	0.4	6.44	< 0.1
	8	8.47	< 0.1	9.61	0.2
20	2	0.05	0.1	3.34	3.1
	5	2.57	1.1	7.97	1.3
	8	5.11	0.2	15	1.3
50	2	0.49	2.3	1.74	4.2
	5	0.54	5	8.2	13.5
	8	1.62	6.8	12.42	33.4
100	2	0.08	11.1	3.32	22.8
	5	1.5	13.6	10.75	40.9
	8	1.86	11	14.33	47.3
<i>Global average</i>		1.66	3.44	6.39	10.53

the average percentage deviation would be higher. Another explanation can also be considered: the number of processors are fixed in ‘Type-2’ problems, that is $m_i = 5$, and the scheduling problem becomes more difficult to solve for CDADS.

Results show the behavior of our approach with variations of n and m . For a given n , the average percentage deviation increases with increasing m . Indeed, the problem difficulty increases when m increases and the obtained solution is further away from the lower bound. On the other hand, for a given number of stages m , increasing n has no significant effect on the average percentage deviation, as the effectiveness of CDADS is independent of the number of jobs: the stability of our method seems to be not linked to the number of jobs n , since for a given m (e.g., $m = 8$), in ‘Type-1’ problems, when n increases from 50 jobs to 100 jobs, the average percentage deviation increase slightly (from 1.62% to 1.86%). It also can be noticed, that in some cases, increasing n results in a decrease in the deviation value (for the configuration $n = 20, m = 8$ the $\%dev$ is taken to be 5.11%, and is evaluated to 1.62% for $n = 50, m = 8$). Apparently, the lower bound becomes more effective with n increasing.

From the experimental studies, it can be observed that CDADS converges quickly. The average CPU time varies between less than 0.1 seconds and 47.3 seconds. The computational cost is more important in ‘Type-2’ instances, confirming the difficulty of these problems. Similarly, for a fixed m , increasing n

leads to CPU time increase. Conversely, when n is fixed, increasing m increases the CPU time.

5.4 Comparison of CDADS Solutions with State-Of-the-Art Results

Table 3 presents the results of CDADS on $\%dev$, the average percentage deviation (as well as a synthesis of the average CPU time for all instances, in the last line of the table). Furthermore, it shows the results obtained by Jouglet *et al.* in [10]. These results are the most recent and the best-known solutions in literature. Thus, we have compared the results of CDADS with GA (genetic algorithm), CP (constraint programming), and MA (a memetic algorithm that combines GA and CP). We disregard the results published by Ercan *et al.* [14] given inconsistency encountered. We contrast our results only versus those presented in [10]. However, we omit the average deviation published in this latest paper due to detected miscalculation (induced by Ercan *et al.*'s errors). Hence, we recalculated the average percentage deviation for all methods given in [10]. The maximum CPU time is fixed at 900 seconds for GA, CP, and MA.

Table 3. Comparing average percentage deviation (and CPU time)

n	m	‘Type-1’ Problems				‘Type-2’ Problems			
		<i>CDADS</i>	<i>GA</i>	<i>CP</i>	<i>MA</i>	<i>CDADS</i>	<i>GA</i>	<i>CP</i>	<i>MA</i>
5	2	0	0.29	0	0	0	1.23	0	0
	5	0.21	1.35	0	0	0.46	1.44	0	0
	8	1.71	4.15	0	0	0.5	2.38	0	0
10	2	0	0	0	0	1.72	2.83	1.72	1.75
	5	0.66	1.64	0	0	6.44	7.8	6.1	5.67
	8	8.47	9.38	10.32	8.02	9.61	10.87	8.37	8.8
20	2	0.05	0.44	2.59	0.66	3.34	3.7	6.72	3.43
	5	2.57	3.49	10.85	2.78	7.97	9.57	22.86	9.57
	8	5.11	5.69	17.98	5.32	15	17.26	28.52	16.02
50	2	0.49	0.63	2.79	0.49	1.74	2.76	6.54	2.21
	5	0.54	0.59	5.3	0.51	8.2	10.95	20.01	10.32
	8	1.62	2.17	14.42	1.71	12.42	15.89	30.06	17.25
100	2	0.08	0.15	1.96	0.07	3.32	3.05	5.68	2.7
	5	1.5	2.5	5.19	2.33	10.75	14.95	19.13	14.37
	8	1.86	1.99	9.47	2.15	14.33	20.06	23.15	17.83
<i>Global average</i>		1.66	2.27	5.39	1.6	6.39	7.28	11.92	8.32
<i>Average CPU(s)</i>		3.44	879.93	320.3	326.01	10.53	879.08	423.09	511.27

As revealed in Table 3 (and as already noticed in Table 2), on the whole, the total average of $\%dev$ obtained by CDADS is 1.66% and 6.39% for the ‘Type-1’

and ‘Type-2’ problems, respectively. Compared to the corresponding averages of 2.27% and 7.28% achieved by GA, and the corresponding values of 5.39% and 11.92 % obtained by CP, CDADS outperforms the GA and CP algorithms. Furthermore, CDADS was clearly superior to CP especially for larger instances ($n = 50$ and $n = 100$).

As depicted in the table, MA finds slightly better solutions in ‘Type-1’ problems, that is 1.60% is obtained by MA while CDADS gives an average deviation percentage of 1.66%. Overall, CDADS outperforms significantly MA, as CDADS results are at 6.39% from optimal solutions (or lower bounds) for ‘Type-2’ problems against 8.32% for MA.

To further assess the effectiveness of CDADS, we measure the number of improved known solutions. It can be seen from Table 4 that CDADS improves 75 known solutions among the 300 tested instances. Thus, the rate of improvement reaches 25%. The results also outline that most improvements are spotted in large instances ($n = 50, 100$), see figure 2. No significant improvements are noticed for small instances ($n = 5, 10$) since all optimal solutions for these problems are known.

Table 4. Number of improved solutions

n	‘Type-1’ Problems	‘Type-2’ Problems
5	0	0
10	1	0
20	5	10
50	8	20
100	8	23
total	22	53

In this study, we also compare the convergence of algorithms. It can be seen from the last line of Table 3, that CDADS outperforms the genetic algorithm (GA), constraint programming (CP), and the memetic algorithm (MA). Indeed, CDADS takes between less than 0.1 seconds (for small problems) and 47.3 seconds (for large problems) to find their solutions, while methods proposed in [10] converge much more slower [0.7 sec, 900 sec]. Even all results were obtained under different computational budgets, we can conclude that CDADS demonstrates fast convergence. Indeed, according to Dongarra’s normalized coefficients [7], our machine is approximately only 3.5 times faster than the machine used by Jouglet *et al.*




Fig. 2. Variation of the number of improved solutions with the number of jobs

6 Conclusions

In this paper, the hybrid flow shop problem with multiprocessor tasks is addressed by means of a discrepancy search method. The proposed method, Climbing Depth-bounded Adjacent Discrepancy Search (CDADS), is based on adjacent discrepancies. We selected several heuristics to generate the initial solution. A lower bound is also proposed to lead a more efficient search. Compared to the best-known results in the literature, CDADS provides better solutions in little CPU time.

In the short-term, we prospect to apply CDADS to simpler problems like classical hybrid flow shop ($size_{ij} = 1, \forall i, j$), widely studied in the literature. Another expected aim would be to adapt the proposed implementation of discrepancy search to more general scheduling problems, in particular the Resource-Constrained Project Scheduling Problem, which still remains one of the most challenging problems in large-scale scheduling.

References

1. J. C. Beck and L. Perron. Discrepancy-bounded depth first search. In *Proceedings of CPAIOR 2000*, pages 8–10, 2000.
2. A. Ben Hmida, M. Haouari, M.-J. Huguet, and P. Lopez. Solving two-stage hybrid flow shop using climbing depth-bounded discrepancy search. *Computers and Industrial Engineering*, 2010. In Press.
3. A. Ben Hmida, M.-J. Huguet, P. Lopez, and M. Haouari. Climbing depth-bounded discrepancy search for solving hybrid flow shop scheduling problems. *European Journal of Industrial Engineering*, 1(2):223–243, 2007.

4. S. Bertel and J.-C. Billaut. A genetic algorithm for an industrial multiprocessor flow shop scheduling problem with recirculation. *European Journal of Operational Research*, 159(3):651–662, 2004.
5. G. Brooks and C. White. An algorithm for finding optimal or near optimal solutions to the production scheduling problem. *Journal of Industrial Engineering*, 16:34–40, 1965.
6. J. Chen and C.-Y. Lee. General multiprocessor task scheduling. *Naval Research Logistics*, 46:57–74, 1999.
7. J. Dongarra. Performance of various computers using standard linear equations software. Technical report, University of Tennessee, 2009.
8. M. F. Ercan and Y.-F. Fung. Real-time image interpretation on a multi-layer architecture. In *Proceedings of IEEE TENCON'99*, pages 1303–1306, 1999.
9. W. D. Harvey and M. L. Ginsberg. Limited discrepancy search. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, volume 1, pages 607–615, Montréal, Québec, Canada, August 1995.
10. A. Jouglet, C. Oğuz, and M. Sevaux. Hybrid flow-shop: a memetic algorithm using constraint-based scheduling for efficient search. *Journal of Mathematical Modelling and Algorithms*, 8:271–292, 2009.
11. J.E. Jr Kelley. The critical-path method: Resources planning and scheduling. In Thompson G.L. and Muth J.F., editors, *Industrial Scheduling*, pages 347–365. Prentice-Hall, Englewood Cliffs, 1963.
12. R. E. Korf. Improved limited discrepancy search. In *Proceedings of the 13th National Conference on Artificial Intelligence (AAAI-96)*, volume 1, pages 286–291, Portland, OR, August 1996.
13. M. Milano and A. Roli. On the relation between complete and incomplete search: an informal discussion. In *Proceedings of CPAIOR 2002*, pages 237–250, Le Croisic, France, 2002.
14. C. Oğuz and M. F. Ercan. A genetic algorithm for hybrid flow-shop scheduling with multiprocessor tasks. *Journal of Scheduling*, 8:323–351, 2005.
15. C. Oğuz, Y.-F. Fung, M. F. Ercan, and X.-T. Qi. Parallel genetic algorithm for a flow shop problem with multiprocessor tasks. In *International Conference on Computational Science*, pages 548–559, Berlin Heidelberg, 2003.
16. C. Oğuz, Y. Zinder, V. Ha Do, A. Janiak, and M. Lichtenstein. Hybrid flow shop scheduling problems with multiprocessor task systems. *European Journal of Operational Research*, 152:115–133, 2004.
17. P. Prosser and C. Unsworth. LDS: testing the hypothesis. Technical Report DCS TR-2008-273, Dept of Computing Science, University of Glasgow, 2008.
18. F. S. Şerifoğlu and G. Ulusoy. Multiprocessor task scheduling in multistage hybrid flow-shops: A genetic algorithm approach. *European Journal of Operational Research*, 55(5):504–512, May 2004.
19. F. S. Şerifoğlu and G. Ulusoy. Multiprocessor task scheduling in multistage hybrid flow-shops: An ant colony system approach. *International Journal of Production Research*, 44(16):3161–3177, 2006.
20. A. Sprecher, R. Kolisch, and A. Drexel. Semi-active, active, and non-delay schedules for the resource-constrained project scheduling problem. *European Journal of Operational Research*, 80(1):94–102, 1995.
21. T. Walsh. Depth-bounded discrepancy search. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI-97)*, volume 2, pages 1388–1395, Nagoya, Japan, August 1997.